

WHAT IS CLAIMED IS:

1                   1.       In a storage environment comprising a plurality of servers, the plurality  
2 of servers including a first server having a file system storing a plurality of files restored from  
3 a backup medium, the plurality of files including one or more data files and one or more tag  
4 files corresponding to data files that have migrated from the file system, a computer-  
5 implemented method of maintaining consistency of the file system of the first server, the  
6 method comprising:

7                   providing first information comprising information related to the plurality of  
8 files stored in the file system of the first server, the first information comprising a plurality of  
9 entries, each entry corresponding to a file and storing status information identifying whether  
10 the file is a tag file or a data file, each entry storing attributes information identifying one or  
11 more attributes of the file;

12                   comparing the plurality of files to information included in the first  
13 information;

14                   identifying, based upon the comparing, at least a first inconsistency where  
15 information associated with a first file from the plurality of files is inconsistent with  
16 information in the first information; and

17                   performing a first set of one or more operations to resolve the first  
18 inconsistency.

1                   2.       The method of claim 1 further comprising:

2                   identifying at least a first entry in the first information that stores status  
3 information identifying a file as a tag file and for which there is no corresponding tag file in  
4 the plurality of files; and

5                   performing a second set of one or more operations for the first entry.

1                   3.       The method of claim 2 wherein performing the second set of one or  
2 more operations comprises:

3                   determining, based upon information in the first entry, if a repository file  
4 exists corresponding to the first entry; and

5                   deleting the first entry from the first information upon determining that a  
6 repository file corresponding to the first entry does not exist.

1                   4.     The method of claim 2 wherein performing the second set of one or  
2 more operations comprises:  
3                   determining, based upon information in the first entry, if a repository file  
4 exists corresponding to the first entry; and  
5                   creating a tag file using information in the first entry upon determining that a  
6 repository file corresponding to the first entry exists.

1                   5.     The method of claim 2 wherein the plurality of files comprises one or  
2 more repository files storing migrated data, the method further comprising:  
3                   providing second information including information related to files stored by  
4 the plurality of servers that have been migrated;  
5                   comparing the one or more repository files to information stored in the first  
6 information and the second information;  
7                   identifying at least one inconsistency based upon comparing the one or more  
8 repository files to information in the first information and the second information; and  
9                   performing a set of one or more actions to resolve the at least one  
10 inconsistency.

1                   6.     The method of claim 1 wherein:  
2                   the first file is a tag file representing a data file that has been migrated from  
3 the file system of the first server;  
4                   identifying at least a first inconsistency comprises determining that the first  
5 information does not include an entry for the first file; and  
6                   performing the first set of one or more operations comprises deleting the first  
7 file.

1                   7.     The method of claim 1 wherein:  
2                   the first file is a tag file representing a data file that has been migrated from  
3 the file system of the first server;  
4                   identifying at least a first inconsistency comprises determining that the first  
5 information includes an entry corresponding to the first file storing status information that  
6 identifies the file as a data file; and  
7                   performing the first set of one or more operations comprises:  
8                   deleting the first file; and

9 deleting the entry in the first information corresponding to the first file.

1 8. The method of claim 1 wherein:

2 the first file is a tag file representing a data file that has been migrated from  
3 the file system of the first server;

4 identifying at least a first inconsistency comprises determining that the first  
5 information includes an entry corresponding to the first file storing status information that  
6 identifies the file as a tag file, wherein the attributes information stored by the entry does not  
7 match at least one attribute of the first file; and

8 performing the first set of one or more operations comprises updating the first  
9 file such that the at least one attribute of the first file matches the attributes information stored  
10 in the entry in the first information corresponding to the first file.

1 9. The method of claim 1 wherein:

2 the first file is a data file;

3 identifying at least a first inconsistency comprises determining that the first  
4 information does not include an entry for the first file; and

5 performing a first set of one or more operations comprises adding an entry to  
6 the first information for the first file.

1 10. The method of claim 1 wherein:

2 the first file is a data file;

3 identifying at least a first inconsistency comprises determining that the first  
4 information includes a first entry corresponding to the first file storing status information that  
5 identifies the file as a tag file; and

6 performing the first set of one or more operations comprises:

7 adding a second entry to the first information for the first file; and

8 creating a tag file to corresponding to the information in the first entry  
9 in the first information.

1 11. The method of claim 1 wherein:

2 the first file is a data file;

3 identifying at least a first inconsistency comprises determining that the first  
4 information includes an entry corresponding to the first file storing status information that

identifies the file as a data file, wherein the attributes information stored by the entry does not match at least one attribute of the first file; and

performing the first set of one or more operations comprises updating the information in the first information entry corresponding first file such that the attributes information stored in the entry matches the at least one attribute of the first file.

12. In a storage environment comprising a plurality of servers, the plurality of servers including a first server having a file system storing a plurality of files restored from a backup medium, the plurality of files including one or more data files and one or more tag files corresponding to data files that have migrated from the file system, a computer-implemented method of maintaining consistency of the file system of the first server, the method comprising:

providing first information including information related to files stored in the file system of the first server;

providing second information comprising a plurality of entries, each entry storing information related a file stored by the plurality of servers that has been migrated;

comparing a first tag file from the plurality of files to information stored in the second information;

identifying, based upon the comparing, at least a first inconsistency where information associated with the first tag file is inconsistent with the information included in the second information; and

performing a first set of one or more actions to resolve the first inconsistency.

13. The method of claim 12 wherein:

identifying the first inconsistency comprises determining that the second information does not include an entry for the first tag file; and

performing the first set of one or more actions comprises deleting the first tag file.

14. The method of claim 12 wherein:

identifying the first inconsistency comprises determining that the second information includes a first entry for the first tag file, a repository file exists for the first tag file, and information in the first entry does not match at least one attribute of the first tag file; and

performing the first set of one or more actions comprises:

7                    updating the first tag file such that the at least one attribute of the first  
8 tag file matches the information in the first entry; and  
9                    adding an entry to the first information for the first tag file, the entry  
10 storing information related to the first tag file.

1                    15.    The method of claim 12 wherein:  
2                    identifying the first inconsistency comprises determining that the second  
3 information includes a first entry for the first tag file, a repository file exists for the first tag  
4 file, and information in the first entry matches one or more attributes of the first tag file; and  
5                    performing the first set of one or more actions comprises adding an entry to  
6 the first information for the first tag file, the entry storing information related to the first tag  
7 file.

1                    16.    The method of claim 12 further comprising:  
2                    identifying at least a first entry in the second information for which there is no  
3 corresponding tag file in the plurality of files; and  
4                    performing a second set of one or more operations for the first entry.

1                    17.    The method of claim 16 wherein performing the second set of one or  
2 more operations comprises:  
3                    determining, based upon information in the first entry in the second  
4 information, if a repository file exists corresponding to the first entry; and  
5                    deleting the first entry from the second information upon determining that a  
6 repository file corresponding to the first entry does not exist.

1                    18.    The method of claim 16 wherein performing the second set of one or  
2 more operations comprises:  
3                    determining, based upon information in the first entry, if a repository file  
4 exists corresponding to the first entry; and  
5                    creating a tag file using information in the first entry for the second  
6 information upon determining that a repository file corresponding to the first entry exists.

1                    19.    In a storage environment managed by a hierarchical storage  
2 management application comprising a plurality of servers, the plurality of servers including a  
3 first server having a file system storing a plurality of files including one or more data files

4 and one or more tag files corresponding to data files that have migrated from the file system,  
5 a computer-implemented method of recovering information, the method comprising:

6 providing first information including information related to one or more data  
7 files that have been migrated, wherein the information related to each data file that has been  
8 migrated includes information identifying a server and a volume from which the data file is  
9 migrated, and information identifying a server and volume where the migrated data of the  
10 data file is stored, the first information comprising information related to a first data file that  
11 has been migrated;

12 determining, based upon the first information, that the file system does not  
13 contain a tag file corresponding to first data file; and

14 creating a tag file corresponding to the first data file based upon information  
15 included in the first information.

1 20. In a storage environment comprising a plurality of servers, the plurality  
2 of servers including a first server having a file system storing a plurality of files restored from  
3 a backup medium, the plurality of files including one or more data files and one or more tag  
4 files corresponding to data files that have migrated from the file system, a computer program  
5 product stored on a computer-readable medium for maintaining consistency of the file system  
6 of the first server, the computer program product comprising:

7 code for providing first information comprising information related to the  
8 plurality of files stored in the file system of the first server, the first information comprising a  
9 plurality of entries, each entry corresponding to a file and storing status information  
10 identifying whether the file is a tag file or a data file, each entry storing attributes information  
11 identifying one or more attributes of the file;

12 code for comparing the plurality of files to information included in the first  
13 information;

14 code for identifying, based upon the comparison, at least a first inconsistency  
15 where information associated with a first file from the plurality of files is inconsistent with  
16 information in the first information; and

17 code for performing a first set of one or more operations to resolve the first  
18 inconsistency.

1 21. The computer program product of claim 20 further comprising:

code for identifying at least a first entry in the first information that stores status information identifying a file as a tag file and for which there is no corresponding tag file in the plurality of files; and  
code for performing a second set of one or more operations for the first entry.

22. The computer program product of claim 21 wherein the code for performing the second set of one or more operations comprises:  
code for determining, based upon information in the first entry, if a repository file exists corresponding to the first entry; and  
code for deleting the first entry from the first information upon determining that a repository file corresponding to the first entry does not exist.

23. The computer program product of claim 21 wherein the code for performing the second set of one or more operations comprises:  
code for determining, based upon information in the first entry, if a repository file exists corresponding to the first entry; and  
code for creating a tag file using information in the first entry upon determining that a repository file corresponding to the first entry exists.

24. The computer program product of claim 21 wherein the plurality of files comprises one or more repository files storing migrated data, the computer program product further comprising:  
code for providing second information including information related to files stored by the plurality of servers that have been migrated;  
code for comparing the one or more repository files to information stored in the first information and the second information;  
code for identifying at least one inconsistency based upon comparing the one or more repository files to information in the first information and the second information;  
and  
code for performing a set of one or more actions to resolve the at least one inconsistency.

25. The computer program product of claim 20 wherein:  
the first file is a tag file representing a data file that has been migrated from the file system of the first server;

4           the code for identifying at least a first inconsistency comprises code for  
5 determining that the first information does not include an entry for the first file; and  
6           the code for performing the first set of one or more operations comprises code  
7 for deleting the first file.

1           26.    The computer program product of claim 20 wherein:  
2           the first file is a tag file representing a data file that has been migrated from  
3 the file system of the first server;  
4           the code for identifying at least a first inconsistency comprises code for  
5 determining that the first information includes an entry corresponding to the first file storing  
6 status information that identifies the file as a data file; and  
7           the code for performing the first set of one or more operations comprises:  
8                   code for deleting the first file; and  
9                   code for deleting the entry in the first information corresponding to the  
10 first file.

1           27.    The computer program product of claim 20 wherein:  
2           the first file is a tag file representing a data file that has been migrated from  
3 the file system of the first server;  
4           the code for identifying at least a first inconsistency comprises code for  
5 determining that the first information includes an entry corresponding to the first file storing  
6 status information that identifies the file as a tag file, wherein the attributes information  
7 stored by the entry does not match at least one attribute of the first file; and  
8           the code for performing the first set of one or more operations comprises code  
9 for updating the first file such that the at least one attribute of the first file matches the  
10 attributes information stored in the entry in the first information corresponding to the first  
11 file.

1           28.    The computer program product of claim 20 wherein:  
2           the first file is a data file;  
3           the code for identifying at least a first inconsistency comprises code for  
4 determining that the first information does not include an entry for the first file; and  
5           the code for performing a first set of one or more operations comprises code  
6 for adding an entry to the first information for the first file.



1                   29.     The computer program product of claim 20 wherein:  
2                   the first file is a data file;  
3                   the code for identifying at least a first inconsistency comprises code for  
4 determining that the first information includes a first entry corresponding to the first file  
5 storing status information that identifies the file as a tag file; and  
6                   the code for performing the first set of one or more operations comprises:  
7                   code for adding a second entry to the first information for the first file;  
8 and  
9                   code for creating a tag file to corresponding to the information in the  
10 first entry in the first information.

1                   30.     The computer program product of claim 20 wherein:  
2                   the first file is a data file;  
3                   the code for identifying at least a first inconsistency comprises code for  
4 determining that the first information includes an entry corresponding to the first file storing  
5 status information that identifies the file as a data file, wherein the attributes information  
6 stored by the entry does not match at least one attribute of the first file; and  
7                   the code for performing the first set of one or more operations comprises code  
8 for updating the information in the first information entry corresponding first file such that  
9 the attributes information stored in the entry matches the at least one attribute of the first file.

1                   31.     In a storage environment comprising a plurality of servers, the plurality  
2 of servers including a first server having a file system storing a plurality of files restored from  
3 a backup medium, the plurality of files including one or more data files and one or more tag  
4 files corresponding to data files that have migrated from the file system, a computer program  
5 product stored on a computer-readable medium for maintaining consistency of the file system  
6 of the first server, the computer program product comprising:  
7                   code for providing first information including information related to files  
8 stored in the file system of the first server;  
9                   code for providing second information comprising a plurality of entries, each  
10 entry storing information related a file stored by the plurality of servers that has been  
11 migrated;  
12                   code for comparing a first tag file from the plurality of files to information  
13 stored in the second information;

14 code for identifying, based upon the comparing, at least a first inconsistency  
15 where information associated with the first tag file is inconsistent with the information  
16 included in the second information; and  
17 code for performing a first set of one or more actions to resolve the first  
18 inconsistency.

1 32. The computer program product of claim 31 wherein:  
2 the code for identifying the first inconsistency comprises code for determining  
3 that the second information does not include an entry for the first tag file; and  
4 code for performing the first set of one or more actions comprises code for  
5 deleting the first tag file.

1 33. The computer program product of claim 31 wherein:  
2 the code for identifying the first inconsistency comprises code for determining  
3 that the second information includes a first entry for the first tag file, a repository file exists  
4 for the first tag file, and information in the first entry does not match at least one attribute of  
5 the first tag file; and  
6 the code for performing the first set of one or more actions comprises:  
7 code for updating the first tag file such that the at least one attribute of  
8 the first tag file matches the information in the first entry; and  
9 code for adding an entry to the first information for the first tag file,  
10 the entry storing information related to the first tag file.

1 34. The computer program product of claim 31 wherein:  
2 the code for identifying the first inconsistency comprises code for determining  
3 that the second information includes a first entry for the first tag file, a repository file exists  
4 for the first tag file, and information in the first entry matches one or more attributes of the  
5 first tag file; and  
6 the code for performing the first set of one or more actions comprises code for  
7 adding an entry to the first information for the first tag file, the entry storing information  
8 related to the first tag file.

1 35. The computer program product of claim 31 further comprising:  
2 code for identifying at least a first entry in the second information for which  
3 there is no corresponding tag file in the plurality of files; and

4 code for performing a second set of one or more operations for the first entry.

1 36. The computer program product of claim 35 wherein the code for  
2 performing the second set of one or more operations comprises:

3 code for determining, based upon information in the first entry in the second  
4 information, if a repository file exists corresponding to the first entry; and

5 code for deleting the first entry from the second information upon determining  
6 that a repository file corresponding to the first entry does not exist.

1 37. The computer program product of claim 35 wherein the code for  
2 performing the second set of one or more operations comprises:

3 code for determining, based upon information in the first entry, if a repository  
4 file exists corresponding to the first entry; and

5 code for creating a tag file using information in the first entry for the second  
6 information upon determining that a repository file corresponding to the first entry exists.

1 38. In a hierarchical storage management environment comprising a  
2 plurality of servers, the plurality of servers including a first server having a file system  
3 storing a plurality of files including one or more data files and one or more tag files  
4 corresponding to data files that have migrated from the file system, a computer program  
5 product stored on a computer-readable medium for recovering information, the computer  
6 program product comprising:

7 code for providing first information including information related to one or  
8 more data files that have been migrated, wherein the information related to each data file that  
9 has been migrated includes information identifying a server and a volume from which the  
10 data file is migrated, and information identifying a server and volume where the migrated  
11 data of the data file is stored, the first information comprising information related to a first  
12 data file that has been migrated;

13 code for determining, based upon the first information, that the file system  
14 does not contain a tag file corresponding to first data file; and

15 code for creating a tag file corresponding to the first data file based upon  
16 information included in the first information.

1 39. A storage system managed using a hierarchical storage management  
2 application, the storage system comprising:

3 a first server; and  
4 a set of one or more storage units coupled to the first server, the set of storage  
5 units storing a plurality of files restored from a backup medium, the plurality of files  
6 including one or more data files and one or more tag files corresponding to data files that  
7 have migrated from the set of storage units;  
8 a memory configured to store first information comprising information related  
9 to the plurality of files stored on the set of storage units, the first information comprising a  
10 plurality of entries, each entry corresponding to a file, each entry storing status information  
11 identifying whether the file is a tag file or a data file, each entry storing attributes information  
12 identifying one or more attributes of the file;  
13 wherein the first server is configured to  
14 compare the plurality of files to information included in the first  
15 information,  
16 identify, based upon the comparison, at least a first inconsistency  
17 where information associated with a first file from the plurality of files is inconsistent with  
18 information in the first information, and  
19 perform a first set of one or more operations to resolve the first  
20 inconsistency.

1 40. A storage system managed using a hierarchical storage management  
2 application, the storage system comprising:  
3 a first server; and  
4 a set of one or more storage units coupled to the first server, the set of storage  
5 units storing a plurality of files restored from a backup medium, the plurality of files  
6 including one or more data files and one or more tag files corresponding to data files that  
7 have migrated from the set of storage units; and  
8 a memory configured to store first information and second information, the  
9 first information including information related to files stored by the set of storage units, the  
10 second information storing information for one or more files stored by the set of storage units  
11 that have been migrated;  
12 wherein the first server is configured to  
13 compare a first tag file from the plurality of files to information stored  
14 in the second information,

15                                identify, based upon the comparing, at least a first inconsistency where  
16 information associated with the first tag file is inconsistent with the information included in  
17 the second information, and  
18                                perform a first set of one or more actions to resolve the first  
19 inconsistency.

1                                41.     In a storage system managed by a hierarchical storage management  
2 application, the storage system comprising a plurality of servers, the plurality of servers  
3 including a first server having a file system storing a plurality of files including one or more  
4 data files and one or more tag files corresponding to data files that have migrated from the  
5 file system, a data processing system comprising:  
6                                a processor; and  
7                                a memory configured to store first information including information related  
8 to one or more data files that have been migrated, wherein the information related to each  
9 data file that has been migrated includes information identifying a server and a volume from  
10 which the data file is migrated, and information identifying a server and volume where the  
11 migrated data of the data file is stored, the first information comprising information related to  
12 a first data file that has been migrated, the memory further configured to store a plurality of  
13 instructions which when executed by the processor cause the processor to:  
14                                determine, based upon the first information, that the file system does  
15 not contain a tag file corresponding to first data file; and  
16                                create a tag file corresponding to the first data file based upon  
17 information included in the first information.